REMARKS

In the immediately preceding office action, the Examiner rejected the independent claims under 35 U.S.C. 102(a) as being unpatentable over Sharma (U.S. Patent No. 6,108,737) and under 35 U.S.C. 103(a) as being unpatentable over Hum (U.S. Patent Application Publication No. 2004/0123047) in view of Bauman (U.S. Patent No. 6,189,078). The withdrawal of the prior rejections is gratefully acknowledged.

Claims 1-20 are pending. Independent claims 1, 11, and 20 were rejected under 35 U.S.C. 103(a) as being unpatentable over Hum in view of Sharma.

In the previous office action response, the Applicants submit that Hum does not teach or suggest a cluster of processors interconnected in a point-to-point architecture based on description associated with Figure 7 in Hum. The Examiner noted that Figure 7 is merely one embodiment and argues that Hum does make reference to point-to-point connections. However, the Applicants respectfully submit that Hum does not teach a "cluster of processors interconnected in a point-to-point architecture.

Hum has ample opportunity to describe such a point-to-point architecture. However, Hum instead only focuses lengthy description on a bus based architecture shown in Figure 7 and merely makes reference to some point-to-point connections without even clearly describing specifically what the links are connecting. Of course other embodiments are always contemplated, but Hum only makes vague references to point-to-point connections. The only Figure that possibly shows point-to-point connections is Figure 1, and Figure 1 only shows connections between a cluster and an agent. Figure 1 does not show how the cluster is connected to the agent. Each cluster may include its own bus, arbitrator, and bus bridge. Each cluster would then include multiple processors connected over a bus which is then connected to an agent. Even if each cluster includes only one processor as mentioned as a possibility by the Examiner, the processor would still be connected to a bus and a bus bridge to an agent. This is not a cluster of processors interconnected in a point-to-point architecture as recited in the independent claims.

Hum does make vague references to point-to-point connections. The Examiner states that Hum explicitly teaches that a point-to-point connection exists between nodes (Hum paragraph 57). However, paragraph 57 merely states "the number of point-to-point connections a node requires grows with the number of nodes involved." It is unclear whether this means point-to-point connections between a node and cache, between a cluster and an agent, between agents and other agents, or something else, because all of these would require more point-to-point connections as the number of nodes grows. It is improper to assume that the statement "the number of point-to-point connections required grows with the number of nodes involved" means that Hum describes "a plurality of processors interconnected in a point-to-point architecture." The other two reference to point-to-point connections in Hum are equally vague and could mean point-to-point connections between a variety of different entities. Hum has ample opportunity to more clearly describe its architecture, and when it does, it only describes a bus architecture for connecting multiple processors in Figure 7.

The Examiner also argues that Hum teaches a remote data cache in its description of an import cache 250. The Applicants respectfully disagree. The import cache 250 "can also avoid broadcasting requests from elsewhere in the system to the agent's local cluster. The agent can use the import cache to determine that no nodes within the cluster have a copy of the cache line" (Hum paragraph 65) "In one embodiment, the agent maintains directory information about all cache lines for which Home is inside the cluster, but which have been cached outside the cluster." (Hum paragraph 66) Of course, this is theoretically only one embodiment, but no others are described. "The Export Directory can also maintain a cached copy of line for which the Home node is within its local cluster." (Hum paragraph 67)

By contrast, a remote data cache holds data from remote clusters, not from within the cluster. To facilitate prosecution, independent claim 1 has been amended to recite "wherein the remote data cache receives data and state information for memory lines of remote clusters." This element is supported in page 21, lines 22-33; Figure 8 and associated description; page 8, lines 6-10; Figure 9 and associated description; Abstract. More specifically, "remote data caches receive data and state information for memory lines held in remote clusters" (Abstract). By contrast, the import cache described in Hum only maintains directory information about all cache lines "for which Home is inside the cluster" (Hum paragraph 66) and the export directory only maintains a

cached "copy of line for which the Home node is within its local cluster." (Hum paragraph 67) Consequently, the import cache and the export directory are not a remote data cache as recited in all independent claims and does not "receive data and state information for memory lines held in remote clusters" as recited in independent claim 1.

Furthermore, the Applicants respectfully submit that it would be inappropriate to combine Hum and Sharma because Sharma in fact teaches away from the elements recited in the independent claims. The independent claims recite "providing response information with a completion indicator to the processor when it is determined that the cache access request can be handled locally" upon making a determination using a remote data cache. The Examiner argues that Sharma describes a type 0 commit-signal (column 15, lines 33-59) and argues that a type 0 commit-signal is a completion indicator. Even if we assume a type 0 commit-signal is a completion indicator, Sharma actually teaches away from using a completion indicator in a combination with Hum.

Sharma describes a variety of commit-signals, type 0 – type 3. "A first type of commit-signal is type 0 which corresponds to a local command issued by a source processor. The term "local command" denotes that the memory reference operation address, e.g., address x, and the data x present at that address are located in the memory address space of the source node in which the source processor resides." (column 15, lines 33-38) That is, Sharma teaches that a type 0 commit signal should only be used when the address of the memory line is in the local cluster. When the address of the memory line is not in the local cluster, type 1-3 commit signals are used that indicate that a processor should wait.

By contrast, the claim elements recite sending a completion indicator even if the memory line is not in the local/request cluster, i.e. a remote data cache is used. Sharma suggests that a completion indicator, assuming a type 0 commit signal is a completion indicator, should only be sent when the address of the memory line is in the local cluster and other types of commit signals should be used when this is not the case.

In light of the above remarks relating to independent claims and certain dependent claims, the remaining dependent claims are believed allowable for at least the reasons noted above. Applicants believe that all pending claims are allowable and respectfully request a Notice of Allowance for this application from the Examiner. Should the Examiner believe that a telephone conference would expedite the prosecution of this application, the undersigned can be reached at the telephone number set out below.

Respectfully submitted,
BEYER WEAVER & THOMAS, LLP

Godfrey K. Kwan
Reg. No. 46,850

P.O. Box 70250
Oakland, CA 94612-0250
(510) 663-1100